

Bank and Allocation Manual

Table of Contents

Preface	3
Introduction (Definitions)	4
Comparison of Relevant Tools	7
Comparison Chart	7
Account Tools (Obsolete)	8
Bank Tools	9
Usage Tools	10
Former Usage Tools and LRMUSAGE	12
LRMUSAGE Execute Line	14
LRMUSAGE Options	16
LRMUSAGE Examples	21
User-Oriented Examples	21
Bank-Oriented Examples	24
Division-Oriented Examples	27
Web-Based Accounting Report Generator	29
LRMUSAGE Web Interface	30
HPSS Reports Interface	34
Static Reports Web Interface	35
Managing A Bank	36
BRLIM Resource Limits Reporting Tool	38
BRLIM Execute Line	38
BRLIM Options	40
BRLIM Examples	41
UINFO Bank Membership Reporting Tool	43
Disclaimer	44
Keyword Index	45
Alphabetical List of Keywords	46
Date and Revisions	47

Preface

Scope: This Bank and Allocation Manual is a reference document that explains the (not always obvious) relationships between usage, divisional banks, and computer-resource allocations on LC computers, and the connection of all three to LCRM (formerly DPCS, the LC "batch" system). This manual systematically compares resource-limit, bank, and allocation-usage software, and reveals the role of each software tool. It fully describes the features and options of LC's time-usage reporting utility (LRMUSAGE, formerly called PCSUSAGE), gives sample annotated output for its most typical reports, and introduces the alternative web-based interface to LRMUSAGE. Finally, it introduces the key features of the BRLIM resource-limit reporting tool, which complements LRMUSAGE, as well as the UINFO bank-membership reporting tool.

This manual supplements the LCRM (DPCS) Reference Manual (URL: <http://www.llnl.gov/LCdocs/dpcs>), which explains the internal workings of the LC batch system (including fair-share scheduling), and the EZJOBCONTROL Basic Guide (URL: <http://www.llnl.gov/LCdocs/ezjob>), which tells how to prepare and run batch jobs on LC machines.

Availability: The programs described here support "production computing" at LC and so are available on every LC production machine, open or secure.

Consultant: For help contact the LC customer service and support hotline at 925-422-4531 (open e-mail: lc-hotline@llnl.gov, SCF e-mail: lc-hotline@pop.llnl.gov).

Printing: The print file for this document can be found at:

OCF: <http://www.llnl.gov/LCdocs/banks/banks.pdf>
SCF: https://lc.llnl.gov/LCdocs/banks/banks_scf.pdf

Introduction (Definitions)

LC manages and tracks your use of computer resources using the RAC ("resource allocation and control") database, one of two major parts of LCRM (the Livermore Computing Resource Management system, formerly called DPCS, the Distributed Production Control System). The LCRM (DPCS) Reference Manual (URL: <http://www.llnl.gov/LCdocs/dpcs>) explains this larger scheme and diagrams (URL: <http://www.llnl.gov/LCdocs/dpcs/index.jsp?show=s2>) its overall structure. (Starting in 2003, DPCS changed its name to LCRM, but most user tools remain unchanged.) Users interested in tracking time usage can concentrate, however, as this Bank and Allocation Manual does, on the LCRM software devoted to managing usage, banks, and allocations. This section explains the difference between accounts (now obsolete), banks, and allocations as historical and conceptual background for the rest of the manual, where individual software tools are described.

FAIR-SHARE Background.

All LC production machines now use a so-called "fair-share" approach instead of a traditional fixed-allocation approach to time management. Traditional and fair-share schemes both address the "political problem" of strategically dividing available resources among users (and user groups). But they approach the political problem from quite different perspectives.

Traditional schemes view resource allocation as primarily a *banking* problem. The scheme assigns each user (or bank) some CPU time for which they are held responsible, and it encourages them to spend it wisely (by the amounts replenished, the timing of allocation shifts, etc.).

Fair-share schemes view resource allocation as primarily a *scheduling* problem. By contrast, they distribute no time amounts to users or banks before job execution. Instead, the fair-share approach rewards those who compute "fairly" by scheduling them to continue to consume resources (by running jobs with a higher priority and hence at a greater consumption level than those who compute "less fairly"). At LC, the fair-share scheduling algorithms are designed to:

- Distribute compute resources to competing jobs in proportion to their user's "shares" (entitlements to compute, a measure of your relative importance).
- Balance the system load through time by rewarding users who spread their work relative to others, and
- Avoid the big "time edges" of traditional allocation exhaustion and sudden shift-based replenishment. A user has no time stockpile and so never runs out of time, but instead sees their scheduling priority constantly readjusted to reflect the "fairness" of their past work.

For details on how the fair-share approach is implemented at LC, see the Fair Share Scheduling (URL: <http://www.llnl.gov/LCdocs/dpcs/index.jsp?show=s5>) section of the LCRM (DPCS) Reference Manual. In light of this approach, accounts, banks, allocations, and usage have the following interpretations at LC:

ACCOUNTS.

Historically at LC, every user was assigned one or more recharge accounts (not to be confused with your username or "login account"). These were required names or numbers (once linked to a time credit) used to label and track the computer time consumed by each user's interactive sessions and batch jobs. Hence, until 2005, much software existed (usually betrayed by the strings "acc" or "-a") to create, assign, and change accounts and to associate specific accounts with specific jobs. Such accounts helped reveal the

total computational "effort" expended when several users ran the same production code or collaborated on the same project.

Changes in LC allocation policy, reflected in iterative changes in LCRM software, have gradually reduced to zero the importance of recharge accounts (and therefore the usefulness of LC's former account-management tools). In August, 2005, version 6.13 of LCRM completely eliminated each job's former account attribute and rendered all former account software obsolete (most has been removed from LC production machines). Only an optional 127-character "project name" remains as a nonaccounting identifier for jobs if a user wishes (set with PSUB's -prj option).

BANKS.

At LC, a bank is a pool of unitless, nondecremented "shares" (or entitlements to use computing resources), recorded in the RAC database along with attributes that specify how those shares behave:

(1) Every bank has (optional) parent and child relationships with other banks. This bank hierarchy strongly influences batch-job scheduling because shares are assigned and their effects enforced in layers. To see the complete bank hierarchy that applies to any particular LC production machine (including the hierarchical share assignment), log on to that machine and type:

```
pshare -T root
```

(2) Banks use permissions to differentiate between those users whose jobs can invoke their shares and those who cannot (so a bank is indirectly a set of users allowed to invoke it).

(3) Since computing time is a finite resource, banks still serve as an accounting device to group and track related time charges as user jobs assigned to them run and spend time. Only this third aspect concerns us here (see the LCRM (DPCS) Reference Manual (URL: <http://www.llnl.gov/LCdocs/dpcs>) and the EZJOBCONTROL (URL: <http://www.llnl.gov/LCdocs/ezjob>) basic guide for more on the other two aspects of banks). Starting in February, 2006, LCRM makes no distinction between interactive and batch banks when reporting time charges for a job.

ALLOCATIONS.

The traditional approach regarded an allocation as an amount of computer time (associated with a shift, or a set of users); users decremented their time allocation until it was exhausted and LCRM stopped their jobs. Under the fair-share approach (URL: <http://www.llnl.gov/LCdocs/dpcs/index.jsp?show=s5>), allocated shares (which are neither decremented nor exhausted) give each user a priority that affects how their jobs are scheduled. But even though an allocation is no longer a rigidly enforced time quota, computer time is still scarce and tracking its use by individuals and by related groups (by "banks") remains important for resource planning and job management. To help compare the time expense of different jobs and users, and to monitor your time use over a period (of one or more shifts), LC provides several usage-reporting tools whose features are described later in this manual. Remember that while banks (= user groups) remain an important part of such usage monitoring, accounts no longer have any place in it at all.

Some individual machines have local LCRM-enforced limits on CPU time/job. And some clusters of similar nodes (called "LCRM resource partitions") have global LCRM-enforced limits on allowed node-time per user or per bank (or both). These limits are not strictly the same as time allocations, but they can have the same effect on preventing waiting jobs from being scheduled when the limits are hit. For a comparative look at these time limits and the tools that report them, see the section "Step 3: Plan Your Job Constraints" in the EZJOBCONTROL (URL: <http://www.llnl.gov/LCdocs/ezjob>) guide.

USAGE.

(1) CPU Time Spent. In this manual, and certainly in the variety of reports generated by LRMUSAGE (details and examples below), usage means simply CPU time spent (during a specified period). In the LCRM fair-share job priority formula, however, usage means something else entirely: a decayed historical aggregate of (potentially) several resources over a long time, and normalized over the set of currently active users. Hence there is no simple way for you to compute your fair-share scheduling "usage" from the usage (periodic CPU time) reports described and illustrated in this document (but see (2) below). The LRMUSAGE tool featured here is for detailed, retrospective tracking of time spent (who used how much time, where, and when; indeed, the web interface to LRMUSAGE even offers canned monthly retrospective summaries (page 35)). It is not for predicting how your job(s) will compete against other jobs during future scheduling. For a convenient comparison of user-, job-, and scheduling-oriented "usage" reports, see the Usage Tools (page 10) section below.

(2) Aggregate Resource Units. To help you assess how separate batch jobs that you run contribute to your priority-relevant usage (in aggregate resource units or AGUs), PSTAT (optionally) reports the AGU value for each job. If you invoke PSTAT's -f ("full report") option for a specific job (-n *jid*), the "resources used" field in the PSTAT output reveals that job's current AGU value (e.g., 73373.00). This can yield helpful overt AGU comparisons if you run many jobs with different allocations under different conditions.

Comparison of Relevant Tools

Comparison Chart

Many different software tools, developed at different times, pertain to accounts (now obsolete), banks, and time usage (most are for general users, but a few require coordinator privileges). Some are single-task tools, while others (like LRMMGR, formerly called PCSMGR) are very general tools that include important bank-related suboptions. While these tools overlap in a few areas, they generally fall into three distinct groups as shown in this chart:

<u>Account Tools</u>	<u>Bank Tools</u>	<u>Usage Tools</u>
[most account tools are now obsolete]	lrmngr create bank	lrmusage
	bac	[replaces userusage, bankusage, padusage]
	newbank	pshare
psub -a <i>acc</i>	psub -b <i>bnk</i>	pstat -f -n <i>jid</i>
palter -a <i>acc</i>	palter -b <i>bnk</i>	brlim -b <i>bnk</i>
		phstat

NOTE: if you have coordinator privileges, go to the "[Managing A Bank](#)" (page 36) section for the privileged steps to make a new bank, assign users and shares, and reallocate these as you wish.

Account Tools (Obsolete)

In August, 2005, version 6.13 of LCRM completely eliminated each job's former account attribute and rendered all former LC account-management tools obsolete. Among the now-deleted tools were ACC, NEWACCT, and DEFACCT, as well as the "create account" feature of LRMMGR. To enable a graceful transition, the PSUB and PALTER -a option now temporarily sets the 127-character "project name" (-prj) field that LCRM associates with each batch job.

`psub -a acc` (formerly account name) sets to *acc* the optional 127-character project field associated with the batch job you are using PSUB to submit (other arguments are also needed for a normal submittal; see EZJOBCONTROL). This has no effect on job accounting.

`palter -a acc -n jid` (formerly account name) changes to *acc* the optional 127-character project field associated with the batch job whose PSTAT-revealed job identifier is *jid*. You can only change jobs that you submitted and only while they are queued BEFORE they start to run. This has no effect on job accounting.

Bank Tools

The following tools, shown here with some typical execute lines, manipulate, assign, or report on LC banks. NOTE: if you have coordinator privileges, go to the "[Managing A Bank](#)" (page 36) section for the privileged steps to make a new bank, assign users and shares, and reallocate these as you wish.

`lrmngr create bank bbb`

uses the general LRMMGR (formerly called PCSMGR) utility, which system managers and computer coordinators have permission to run in this way, to create a new bank (resource allocation pool) called *bbb*. Other LRMMGR arguments assign this bank its attributes (such as its parent, its allowed users, or its total shares).

`bac`

reports the parents, children, and permissions (batch or interactive use, scheduling priority, etc.) associated with your current bank or any specified bank (with `-b bnk`). Using `-T bnk` reports on children too, while `-r bnk` reports on parents too. Bank resources and shares are NOT reported.

`newbank bnk`

sets to *bnk* the bank charged for the time used by (all processes within) the session from which NEWBANK is run (otherwise your default bank is charged). NEWBANK -l lists your current bank, while NEWBANK with no arguments lists all valid banks you are authorized to tap (often just one) and prompts you to choose one. The corresponding subroutine to call from LIBPCS (/usr/local/lib/libpcs.a) is PCSSETCURBANK. A DEFBANK program also exists but changing your default bank is seldom appropriate (and for many users not possible). On BlueGene/L, NEWBANK always returns an error (so run it elsewhere).

`psub -b bnk`

sets to *bnk* the bank charged for the batch job you are using PSUB to submit (other arguments are also needed for a normal submittal; see EZJOBCONTROL). Your default (interactive) bank may not be the appropriate bank for charging your batch jobs on all machines, so overtly specifying a bank for each batch job is very desirable in a multimachine environment.

`palter -b bnk -n jid`

changes to *bnk* the bank charged for the batch job whose PSTAT-revealed job identifier is *jid*. You can only change jobs that you submitted and only while they are queued BEFORE they start to run. This is comparable to using NEWBANK interactively.

`uinfo bank bnk`

reports the login name of every user who belongs to the specified *bnk*. UINFO can also report all banks to which any specified user belongs. See the [UINFO section](#) (page 43) below for details.

Usage Tools

The following tools, shown here with some typical execute lines, report "usage" from several different perspectives, including personal or divisional CPU time used, CPU time used by a specific batch job, and fair-share scheduling "usage." They do NOT assign or manipulate bank resources (for that, see the [Bank Tools](#) (page 9) section above).

`lrmusage formatoption scopeoptions`

reports the CPU time in minutes (or another unit if you request) consumed by the users or banks on the machines or over the time range that you specify. You must supply exactly one argumentless *formatoption* that specifies the usage report's primary focus: on banks (-bu), on users (-ub), on machines (-bm, -um), or on whole administrative divisions (-pad, -pad2). With other *scopeoptions* you can select usage information on specific users (-u) or banks (-b), on specific LC production machines (-h) or groups of related machines, or on a specific range of whole days. The default LRMUSAGE report covers just the user who runs it on the current machine for yesterday only. See the [LRMUSAGE](#) (page 14) section below for option details, output restrictions that affect most users, and example reports. Users familiar with the former BANKUSAGE, USERUSAGE, or PADUSAGE tools should also check the [Former Tools](#) (page 12) comparison section below for a discussion of how to mimic those tools by combining appropriate LRMUSAGE options (and how these tools diverge).

`pstat -f -n jid [-T]`

reports the CPU time in minutes (not hours) used so far by the specific batch job whose LCRM-assigned numerical job ID is *jid* (displayed along with about 10 other job properties in a multiline report). If the job has already terminated (within the last 5 days), add the -T option to get a report that includes its total CPU time used. This -f ("full," 34-field) report also includes "resources used," PSTAT's somewhat misleading label for a current estimate of this job's "aggregate resource units" or AGUs, the only usage measure directly relevant to calculating the job's fair-share scheduling priority.

`pshare -p -t bnk -0`

reports the raw and normalized shares, the normalized "decayed usage," and the resulting fair-share priority for all and only the currently active users in bank *bnk* (most relevant if it is your own bank). Usage here is a unitless value between 0 and 1 inclusive that (1) reflects historical CPU time used according to a half-life decay formula and (2) is normalized over (divided by) the sum of decayed usage for each active user. There is no simple way to connect your unitless PSHARE "usage" with your LRMUSAGE "usage" in CPU minutes consumed. See the [Fair Share Scheduling](#) (URL: <http://www.llnl.gov/LCdocs/dpcs/index.jsp?show=s5>) section of the LCRM (DPCS) Reference Manual for the usage algorithms that PSHARE employs.

`brlim -u ulist -b blist`

reports the current partition-wide (not local) job, node, and node-time limits and LCRM resource commitments against these limits for the specified user(s) and bank(s). These commitments against resource limits function somewhat like traditional allocations when LCRM schedules jobs. See the BRLIM section (page 38) below for details and examples. (BRLIM does not run on nor report on LC machines whose jobs are scheduled by Moab.)

`phstat`

reports by machine ("host"), for a list of open (OCF) or secure (SCF) LC hosts, the current job-scheduling and resource commitments for each host (such as total memory load, nodes committed, and backfill availability). These are all *nonbank* features that may affect how (or whether) your job can actually use the bank time that you have planned for it. See the PHSTAT section (URL: <http://www.llnl.gov/LCdocs/ezjob/index.jsp?show=s6.4.2>) of EZJOBCONTROL for details.

Former Usage Tools and LRMUSAGE

In the spring of 2000 a single utility (then called PCSUSAGE, now called LRMUSAGE) replaced three former usage-reporting tools (BANKUSAGE, USERUSAGE, and PADUSAGE). To help users familiar with those old tools find appropriate corresponding features of LRMUSAGE, this section summarizes the former tools and compares their defaults and execute lines to those of LRMUSAGE.

REPLACED TOOLS.

The three replaced tools for reporting "usage" (personal or divisional CPU time actually used) and their roles were:

bankusage	reported CPU time in minutes consumed by every user who charged against a specified bank (where users could control the range of banks reported). LRMUSAGE's -bu and -bm options (with others) simulate these bank reports.
userusage	reported CPU time in minutes consumed by a specified user or set of users (where users could control the time range covered). LRMUSAGE's -ub and -um options (with others) simulate these user reports.
padusage	reported CPU time in minutes consumed collectively by each LLNL administrative division (e.g., B division) during a specified time range (displayed hierarchically by division). LRMUSAGE's -pad and -pad2 options (with others) simulate these divisional reports.

DEFAULT DIFFERENCES.

Mapping between a favorite report generated with the former usage-reporting tools and a comparable report generated by LRMUSAGE is complicated by several major default differences between LRMUSAGE and the tools that it replaced. Some defaults have changed and others must be invoked differently with LRMUSAGE. This table alerts you to the most important default differences:

Output Parameter	Default for LRMUSAGE	Former default for BANKUSAGE, USERUSAGE, PADUSAGE
User range	current user	all users
Bank range	current user's bank(s)	all banks
Time range	yesterday	today
Time units	user setable (-time)	minutes
Time edges	always 0:00 and 24:00, days setable	user setable
Display format	must be specified (-bu, -ub, etc.)	implied by choice of utility
List delimiter	blank space	comma

EXECUTE LINES COMPARED.

The default differences above together with some more subtle changes (in command syntax and in how "partitions" of machines are named, for example) mean that some former execute lines have much shorter replacements with LRMUSAGE while others have longer replacements or no replacement at all. Also, most people who run LRMUSAGE can only get time reports for themselves, for their own bank(s), and for other users who share those bank(s), NOT on any "outside" users or banks (a major policy change from the past). See the sections on LRMUSAGE Options (page 16) and LRMUSAGE Examples (page 21) below for full details on how to construct appropriate LRMUSAGE execute lines to yield specific reports. The typical comparison cases shown here suggest the differences in execution that users of the former tools should watch out for:

Former Execute Line:

bankusage -b ee,me -P smp

userusage -u lflor -n "yesterday"

userusage -u lflor -tb "12:10" -te "13:30"

padusage -P smp [today]

Comparable LRMUSAGE Line:

lrmusage -bu -b ee me -u all -p up_part

lrmusage -ub -u lflor

[not possible with LRMUSAGE]

lrmusage -pad -p up_part -tb jan 03 2007

[an interactive web interface
is also available]

LRMUSAGE Execute Line

LRMUSAGE (formerly called PCSUSAGE) reports how CPU time (in minutes by default) is consumed by user-specified combinations of users (including the current user or all users), banks (including the current user's bank or all banks), and machines (including various groups of related machines). Similar CPU time reports for whole administrative divisions are also available. You can control the time range covered and (many aspects of) the report displayed. By default (run with no arguments) LRMUSAGE reports a brief help file summarizing its command syntax and then ends. There is no interactive prompting, and LRMUSAGE ends after every report that it issues (but a separate [web interface](#) (page 29) is also available).

To run LRMUSAGE on any LC production machine, type

lrusage *options*

where *options* control the format and scope of the time-used report. In fact, the intended execute line for LRMUSAGE is

lrusage *formatoption scopeoptions*

where exactly one format option is followed by as many scope options as needed to control the users, banks, machines, and days covered. The available format and scope options include:

Report Dimension	Control Option(s)
Report format: [must use exactly one]	-bu -bm -ub -um -pad -pad2
Banks covered:	-b[ank] <i>bnklist</i> -T <i>bnklist</i>
Users covered:	-u[ser] <i>ulist</i>
Machines covered: [use no more than one]	-h[ost] <i>hlist</i> -r[esource] <i>rname</i> -p[artition] <i>pname</i> [-class u c uc]
Time range covered:	-tb <i>date1</i> [-te <i>date2</i>]
Time units, threshold:	-time <i>units</i> -m <i>threshold</i>
Display control:	-t[ab] -f <i>field=value...</i>
Help synopsis:	-help [<i>opname</i>]

All list arguments (e.g., lists of host names) are blank (not comma) delimited. LRMUSAGE has many options that start with the same character (e.g., t), so no arbitrary truncation of options is allowed: use the whole option name or its specified abbreviation only (e.g., -h is allowed only for -host but not for -help). You must use exactly one of the report-format options on every LRMUSAGE execute line (unless you want just the help message).

DATE FORMAT:

LRMUSAGE accepts dates (to specify the time range for usage reports) in exactly one format:

mmm dd yyyy

(no surrounding quotes are needed, but allowed), where

<i>mmm</i>	are the first three characters of the desired month's name, in either or mixed case (e.g., nov).
<i>dd</i>	are the digits for the day of the month desired (one or two digits allowed, e.g., 3 or 03).
<i>yyyy</i>	are four digits specifying the year desired (all four digits are required, and the year field can NOT be omitted or defaulted).

The default ending date (-te) is the same as the beginning date (-tb), and the default beginning date is always yesterday, NOT today. To get usage data for today from LRMUSAGE, you MUST overtly invoke -tb with today's date as its argument. Every date starts at 0:00 and ends at 24:00; you can NOT specify a different start time (of day) or end time (of day) in any format.

WARNINGS:

- (1) Requests for reports with 0.00 data (e.g., on yourself for a time range or a machine where you happen to have used no time) often cause LRMUSAGE to fail. The failure message says "no data was found for the following query," followed by the (usually hidden) SQL query that LRMUSAGE tried on your behalf. You can sometimes get such 0.00 reports indirectly, as part of a longer report some of whose entries are nonzero (e.g., try -u all).
- (2) Most people who run LRMUSAGE can only get time reports for themselves, for their own bank(s), and for other users who share those bank(s), NOT on any "outside" users or banks (a major policy change from the past). Requests for data on other users or banks are ignored. Bank managers can see data about any user in any of the (usually hierarchical tree of) banks under their control. Only a few specially authorized bank coordinators and LC staff members can see data on any user or bank.
- (3) A delay of up to 10 minutes can occur between when a job charges time and when LRMUSAGE reports that time as used.

LRMUSAGE Options

Report format (must use exactly one):

- bu creates a bank-oriented report in which time used is organized first by bank and then by user(s) within each bank covered. Time used is summed for each bank covered, and then again as a grand total for all covered banks. You must use other options to specify which banks, users, or time ranges to cover in this format (the default -bu report covers you and your bank(s) on the current machine for yesterday only).
- ub creates a user-oriented report in which time used is organized first by user and then by bank(s) for each user covered. Time used is summed for each user covered, and then again as a grand total for all covered users. You must use other options to specify which banks, users, or time ranges to cover in this format (the default -ub report covers you and your bank(s) on the current machine for yesterday only).
- bm creates a machine-oriented report in which time used is organized first by machine and then by bank(s) and then by user(s) for each machine covered. Time used is itemized by bank and user but is summed only by machine, and then again as a grand total for all covered machines. You must use other options to specify which banks, users, or machines to cover in this format (the default -bm report covers you and your bank(s) on the current machine for yesterday only). This option covers only OCF machines when used on the open network, and only SCF machines when used on the secure network.
- um creates the same machine-oriented report, with the same itemizing, summing, and default scope, as -bm above, except that time used is organized first by machine and then by user(s) and then by bank(s) for each machine covered. This option covers only OCF machines when used on the open network, and only SCF machines when used on the secure network.
- pad creates a hierarchical report by LLNL program and then by administrative division that shows the time used and percentage of total CPU time for each organizational unit. You must use other options (-bank and -user NOT allowed) to specify which machines and time range to cover in this format (the default -pad report covers the current machine for yesterday only).
- pad2 creates an alphabetized report by LLNL division that shows the time used and percentage of total CPU time for each organizational unit. You must use other options (-bank and -user NOT allowed) to specify which machines and time range to cover in this format (the default -pad2 report covers the current machine for yesterday only).

Banks covered (use only one):

-b[ank] *bnklist*

requests that the usage report cover all and only the banks specified in the blank-delimited *bnklist* (default is only the current user's bank(s)). Use -b all to select all banks (a long list).

WARNING: for authorized users only. Most users can only see reports on their own bank(s) regardless of which banks they request with -b.

-T *bnklist*

(total bank) requests that the usage report cover all and only the parent banks in the blank-delimited *bnklist*, as well as all their child banks (default is -b format above, no children).

Users covered:

-u[ser] *ulist*

requests that the usage report cover all and only those users whose login names are specified in the blank-delimited *ulist* (default is only the current user who runs LRMUSAGE). Use -u all to select all users (a long list). The order of login names in *ulist* has no effect on their order in the usage report (which is always alphabetical). LRMUSAGE does NOT reveal the real name associated with the login name(s) that you specify.

WARNING: for authorized users only. Most users can only see reports on themselves and the users in their own bank(s) regardless of which users they request with -u.

Machines covered (use only one):

-h[ost] *hlist*

requests that the usage report cover all and only those computers specified in the blank-delimited *hlist* (default is only the computer on which you run LRMUSAGE). You can include dissimilar computers in *hlist* (e.g., -h west blue). If all the computers on which you want a report belong to the same "resource" or the same "bank partition," you may find it easier to use the -r or -p options (described below).

-r[esource] *rname*

requests a usage report covering all computers in the same "resource," a predefined group of similar machines often managed as a cluster. Using -help r reveals the resource names, which currently include over 30 different machine groups in random order (some physical, some political, and some legacy names for convenience).

-p[artition] *pname* [-class u|c|uc]

requests a usage report covering all computers in the same bank partition, a group of machines managed together by LCRM for accounting purposes. Using -help p reveals the partition names, which currently include over two dozen designated partitions (some physical, some political, and some legacy names for convenience).

Some partitions could be classified or unclassified (such as `asci_part` on both the open and secure networks). So a `-class` suboption exists to specify classified (c), unclassified (u), or dual (uc) report coverage in those cases only.

Time-range covered:

`-tb date1 [-te date2]`

requests that a usage report cover only the entire period starting on *date1* and ending on *date2* (default is yesterday, NOT today), where dates can be specified in exactly one format

mmm dd yyyy

(no surrounding quotes are needed, but allowed), where

mmm are the first three characters of the desired month's name, in either or mixed case (e.g., nov).

dd are the digits for the day of the month desired (one or two digits allowed, e.g., 3 or 03).

yyyy are four digits specifying the year desired (all four digits are required, and the year field can NOT be omitted or defaulted).

`-time units` specifies the unit of time for a LRMUSAGE report, where the choices for *units* are minutes [the default]

hours

days

seconds

If you use `-m` along with `-time`, then your choice of *units* applies to the report threshold as well as to the report output (e.g., both could be hours instead of minutes). You cannot truncate the `-time` option.

`-m threshold` specifies an integer *threshold* value so that LRMUSAGE reports only usage times that exceed this value. The default *threshold* is 0; the default unit for *threshold* is minutes, which you can change by using the `-time` option along with `-m`.

Display control:

`-f field=value,...`

(must be used alone, not with any other options) allows you to construct a query (a comma-delimited list of *field=value* pairs) and pass it to ORACLE to return a customized report on the normally hidden "daily usage table." Here *field* is the name of a column in that table, where the choices include:

auser [user login name]

bank

daily_ucpu [user CPU time]

daily_scpu [system CPU time]

daily_icpu [idle CPU time]

daily_memint [memory integral]

daily_vmemint [virtual memory integral]

lcrm_part [bank partition]

host

jtype [job type]

If you enter a faulty column name, a help message returns the current list of valid column names.

`-t[ab]` returns output in a spare, tab-delimited format (for input to other programs for postprocessing). The default is a labeled, easy to read, space-delimited format.

`-help [option]` lists all currently available LRMUSAGE options (briefly, as a MAN-page synopsis), or, if you include the name of an option, briefly describes that option.

LRMUSAGE Examples

This section shows typical user-oriented, bank-oriented, and then divisional-oriented reports on time used, and a few of the most helpful LRMUSAGE control options in play.

User-Oriented Examples

[1] User Default Report

GOAL: To discover the time used by you only (here with login name ghi), yesterday only (here, January 30), on the current machine only (here, THUNDER). This is LRMUSAGE's default user-oriented report.

STRATEGY: Run LRMUSAGE with the -ub option (and no others). This requests the default report in "user" format (sorted first by user and then by bank). Note that the report heading reveals the default host (cluster name only), time range (always the day before the day that you run LRMUSAGE), reporting threshold, and units. LRMUSAGE may take several *minutes* (wall-clock time) to generate a response.

```
User: lrmusage -ub
Rtne:
```

```
*****
*                LRMUSAGE                *
*****
Report for      : thunder
Start Date      : Jan 30 2007
End Date        : Jan 30 2007
Minimum time    : 0.0 minutes
Time units      : minutes
*****
```

```
User: ghi
Bank                Time Used
-----
cs                  44.88
                   -----
                   44.88
                   -----
                   -----
                   44.88 total
```

[2] Nondefault Time Range

GOAL: To discover the time used by you (login ghi) on the current machine (THUNDER) for a specified nondefault time range (here, from 0:00 on January 1 through 24:00 on January 20). A special case of this is when your intended time range is "today" (in contrast to yesterday, always LRMUSAGE's default time range).

STRATEGY: Run LRMUSAGE with the -ub option ("user" format) as well as -tb to specify the start of the time range and -te to specify its end (use dates only, as reflected in the report heading; you cannot set the time of day to start or stop, and you cannot omit the year). To specify "today," use -tb with today's date as its argument (-te defaults to the same date as -tb). LRMUSAGE may take several *minutes* (wall-clock time) to generate a response to this complex query.

```
User: lrmusage -ub -tb jan 01 2007 -te jan 20 2007
```

```
Rtne:
```

```
*****
*                LRMUSAGE                *
*****
Report for      : thunder
Start Date      : Jan 01 2007
End Date        : Jan 20 2007
Minimum time    : 0.0 minutes
Time units      : minutes
*****
```

```
User: ghi
```

Bank	Time Used
-----	-----
cs	114.67

	114.67

	114.67 total

GOAL: To discover the time used by several specified users on several specified machines (perhaps all different from the machine on which you run LRMUSAGE).

STRATEGY: Run LRMUSAGE with the -ub option ("user" format), the -u option to specify the (login names of the) target users (here ghi and pqr), and the -h option to specify the target hosts (here THUNDER and UP, blank delimited). Yesterday is the default time range. The report heading reflects your choice of hosts, while the report body reflects your choice of users.

WARNING: Security precautions now limit you to report on only those users who draw against the same bank(s) as you (other names are ignored).

```
User: lrmusage -ub -u ghi pqr -h thunder up
```

```
Rtne:
```

```
*****
*                LRMUSAGE                *
*****
Report for      : thunder,up
Start Date     : Feb 02 2007
End Date       : Feb 02 2007
Minimum time   : 0.0 minutes
Time units     : minutes
*****
```

```
User: ghi
```

Bank	Time Used
-----	-----
cs	0.05

	0.05

```
User: pqr
```

Bank	Time Used
-----	-----
cs	4.10

	4.10

	4.15 total

Bank-Oriented Examples

[1] Bank Default Report

GOAL: To discover the time used by you only (here with login name ghi), yesterday only (here, February 2), on the current machine only (here, THUNDER). This is LRMUSAGE's default bank-oriented report (same content as the default user-oriented report but organized by bank). Starting in February, 2006, LCRM does not distinguish between interactive and batch banks when reporting time changes for users or jobs.

STRATEGY: Run LRMUSAGE with the -bu option (and no others). This requests the default report in "bank" format (sorted first by bank and then by user). Note that the report heading reveals the default host (cluster name only), time range (always the day before the day that you run LRMUSAGE), reporting threshold, and units.

```
User: lrmusage -bu
Rtne:
*****
*                LRMUSAGE                *
*****
Report for      : thunder
Start Date      : Feb 02 2007
End Date        : Feb 02 2007
Minimum time    : 0.0 minutes
Time units      : minutes
*****
```

```
Bank: cs
User                Time Used
-----
ghi                  44.88
-----
                     44.88
```

[2] Multiple Banks, Multiple Users

GOAL: To discover the time used by all those whose jobs draw against several specified banks (for yesterday only and the current machine only, the defaults).

STRATEGY: Run LRMUSAGE with the -bu option ("bank" format), the -b option to specify the target banks (here cs and ftpd, blank delimited), and the -u option with the special "all" argument (to select all users in those banks).

WARNING: Most users can only display data on the banks they themselves draw against, and only for other users who also draw against those same banks. Requests for other users or other banks are ignored.

```
User: lrmusage -bu -b cs ftpd -u all
```

```
Rtne:
```

```
*****
*                               *
*               LRMUSAGE        *
*                               *
*****
Report for      : thunder
Start Date      : Feb 02 2007
End Date        : Feb 02 2007
Minimum time    : 0.0 minutes
Time units      : minutes
*****
```

```
Bank: cs
```

User	Time Used
cch	0.00
ghi	44.88
jmw	3.58
kevin	0.02
richmond	0.98
sandra	0.03
taf	0.05

	49.54

```
Bank: ftpd
```

User	Time Used
aro	0.05
awcook	0.00
bolstad	0.02
cah	0.00
chase	0.07
cloutman	0.00
colvin	0.00
dwinter	0.03
hiejima	0.52
kubota	0.00
medwards	0.00
park	0.02
rodrigue	0.00

root	0.00
taylor	0.02
taylorj	0.07
woodin	0.00

	0.80

	50.34 total

Division-Oriented Examples

[1] Hierarchical Division Report

- GOAL:** To discover the time used by all LC (LCRM) administrative divisions, organized and itemized hierarchically by each division and its subdivisions (for yesterday only on the current machine only, the defaults).
- STRATEGY:** Run LRMUSAGE with the -pad option (and no others). Note that the report heading summarizes the time range and machines covered, for reliable later use of this report. LRMUSAGE may take several *minutes* (wall-clock time) to generate a response.

User: lrmusage -pad
Rtne:

```
*****
*                LRMUSAGE                *
*****
Report for : thunder 4024 cpus->Jan 31 2007
Start Date : Jan 31 2007
End Date   : Jan 31 2007
Minimum time : 0.0 minutes
Time units  : minutes
Reference   : 3219840
*****
```

Organization	Time Used	% Delivered
-----	-----	-----
Uncharged Time	355590.7	11.0
- System (kernel; daemons; root processes)	47348.3	1.5
- LC Admin (time used by LC staff)	12125.4	0.4
- Idle (system idle)	266685.0	8.3
- Down Time (system and/or LCRM down)	3371.1	0.1
 Weapons Program	534353.0	16.6
- B-Division	1728.3	0.1
- KX-Division	4365.5	0.1
- WN-Division	0.9	0.0
- Nuclear Design	528258.2	16.4
-- B-Division	6.1	0.0
 Multi-Programmatic Institutional Comp	2839120.5	88.2
- Defense Sciences	534353.0	16.6
 Time reported by LCRM	3194734.1	99.2
- Production (UC+WP+MIC+DAT)	2839120.5	88.2
- Used (Prod+System+Admin)	2898594.2	90.0

[2] Alphabetical Division Report

- GOAL:** To discover the time used by all LC (including ASC) administrative divisions, organized alphabetically by division name (for yesterday only on the current machine only, the defaults).
- STRATEGY:** Run LRMUSAGE with the -pad2 option (and no others). Note that the report heading summarizes the time range and machines covered, for reliable later use of this report. LRMUSAGE may take several *minutes* (wall-clock time) to generate a response.

User: lrmusage -pad2

Rtne:

* LRMUSAGE *

Report for : thunder 4024 cpus->Jan 31 2007

Start Date : Jan 31 2007

End Date : Jan 31 2007

Minimum time : 0.0 minutes

Time units : minutes

Reference : 3219840

Division	Reference	Time Used	Percent
ASCI Turbulence	3219840	0.0	0.0
AX Division	3219840	359310.2	11.2
Alliances	3219840	0.0	0.0
B Division	3219840	160784.0	5.0
CASC	3219840	76184.2	2.4
Chem & Mat Sci	3219840	58543.4	1.8
Down Time	3219840	3371.1	0.1
Energy & Environment	3219840	38363.1	1.2
Engineering	3219840	1776.4	0.1
Homeland Security	3219840	78292.8	2.4
Idle	3219840	266685.0	8.3
Institutional Comput	3219840	656236.4	20.4
LC Admin	3219840	12125.4	0.4
Physics	3219840	178581.2	5.5
System	3219840	47348.3	1.5
Unassigned Users	3219840	1257131.2	39.0
WN Division	3219840	0.9	0.0
		3194734.16	99.3

Web-Based Accounting Report Generator

LC offers a web-based interface to LRMUSAGE (and related software) on both the open and the secure networks at these URLs:

Open: <https://lc.llnl.gov/dfs/www/lcrm/accounting/index.html>

Secure: <https://lc.llnl.gov/dfs/www/lcrm/accounting/index.html>

This interface, which you can access with any HTML browser that supports both DCE authentication (you will be asked for your one-time password (OCF) or DCE password (SCF) every time you visit these sites) and graphical display, offers basically the same functions as does the LRMUSAGE (page 14) utility, with a few options added to plot some results as colored bar or pie charts.

The advantages of the web-based interface include:

- Visual display (as menus) of the choices available as arguments for each LRMUSAGE option,
- Alphabetical listing of options when there are many, and
- Automatic signalling of each network's limitations (only OCF machines are offered, for example, if you use the open network URL).

The disadvantages of the web-based interface include:

- The field descriptions on the web interface often differ markedly from the corresponding command-line option with the identical function (for example, "start date" versus option -tb).
- You can make a script to automatically invoke the same LRMUSAGE options with just the same arguments on repeated runs, but this is not possible with the web interface.
- You must supply your one-time password (OCF) or DCE password (SCF) every time you invoke the web interface, although LRMUSAGE itself has no such requirement.

When you first arrive at either URL listed above, you receive a greeting page titled "Welcome to the ICC Accounting Report Generator." Most of the page is a sample usage bar graph, but along the left-hand edge is a menu of five labeled buttons that invoke this web site's interactive features as follows:

<u>LRMUSAGE</u>	displays a multipart menu of LRMUSAGE options that you can use to specify the accounting report that you want, then have it shown or plotted.
<u>REPORTS</u>	displays a menu of preprocessed (mostly divisional) usage and funding reports (monthly summaries) that you can select for display.
<u>HPSS</u>	(OCF only) offers special reports for authorized users only (not available for general use).
<u>HELP</u>	offers a few sparse menus that outline the report types and output options that this web interface provides.
<u>FEEDBACK</u>	(OCF only) sends an e-mail message to the web-interface programmer (not to the LC Hotline).

LRMUSAGE Web Interface

The web interface of LRMUSAGE menus looks like this (details vary with current machines, and some defunct machines or partitions, like BLUE or GPS_PART, appear even in the current menus):

Select a Report Type			Percent Allocation Delivered	
			Percent by Division	
			Bank by User (time)	
			Bank by User (percent)	
			Bank by Machine (time)	
			Bank by Machine (percent)	
			User by Bank (time)	
			User by Bank (percent)	
			User by Machine (time)	
			User by Machine (percent)	
Selection Options				
	Hosts		Users	Banks
Machine	Partition	Resource		
alc	alc_part	alc		
blue	gps_part	asci_blue	xxx	yyy
frost	ilx_part	gps	all	zzz
gps	penguin_part	purple		aggregate
purple	purple_part	thunder		all
thunder	thunder_part	ubgl		
[others]	[others]	[others]		
Select Time Interval				
Start Date [Jan 15 2007]				
End Date [Jan 15 2007]				
View Calendar				
Output Options				
Format	Tab	Plot	Minutes	0.0
Submit	Reset			

These menus offer the following interactive features (whose corresponding LRMUSAGE commands appear in parentheses):

Select a Report Type

lets you pick one general report format from a mutually exclusive list of alternative formats appearing on a pop-up menu when you click on this field's right-hand side. The alternative report formats are:

Percent Allocation Delivered

(-pad) creates a hierarchical report by LLNL program and administrative division showing the time used and percent of total CPU time used by each division.

Percent By Division	(-pad2) creates a alphabetized report by LLNL program and administrative division showing the time used and percent of total CPU time used by each division.
Bank by User	(-bu) creates a bank-oriented report in which usage is organized first by bank and then by user(s) within each bank covered. You can choose between a report of (1) actual time used or (2) percentage of available CPU time used.
Bank by Machine	(-bm) creates a bank-oriented report in which usage is organized first by bank and then by machine(s) within each bank covered. You can choose between a report of (1) actual time used or (2) percentage of available CPU time used.
User by Bank	(-ub) creates a user-oriented report in which usage is organized first by user and then by bank(s) for each user covered. You can choose between a report of (1) actual time used or (2) percentage of available CPU time used.
User by Machine	(-um) creates a machine-oriented report in which usage is organized first by machine and then by user(s) and bank(s) for each machine covered. You can choose between a report of (1) actual time used or (2) percentage of available CPU time used.

Hosts offers in the three subcolumns under this heading three mutually exclusive ways to specify the machines covered by your accounting report. The alternatives include:

Machine	(-h) requests that the report cover all and only those computers that you pick from the alphabetized, scrollable list of names shown (only OCF names appear on the open-side web site; only SCF names appear on the secure-side web site).
Partition	(-p) requests that the report cover all and only those computers that you pick from the alphabetized, scrollable list of bank partitions shown (only OCF bank partitions appear on the open-side web site; only SCF bank partitions appear on the secure-side web site).
Resource	(-r) requests that the report cover all and only those computers that you pick from the alphabetized, scrollable list of "resource" names shown (only OCF "resource" names appear on the open-side web site; only SCF "resource" names appear on the secure-side web site). For practical purposes, a "resource" is a cluster of like computers.

- Users (-u) specifies the users (by login name) whose usage this report covers. For most people, the only choices are your own login name (shown as xxx in the diagram) and all users (a long list).
- Banks (-b) specifies the banks (by abbreviation) whose usage this report covers. For most people, the only choices are your own bank(s) (shown as yyy and zzz in the diagram) and all banks (a long list). The "aggregate" choice (incompatible with "all" but compatible with any specific banks) adds the individual usage numbers to a total.

Select Time Interval

specifies the time period that your usage report covers (in whole days only). The default is always yesterday, NOT today (which you must explicitly request).

- Start Date (-tb) lets you type into the accompanying text-input field your desired start date. Always use the format shown, including a full four-digit year.
- End Date (-te) lets you type into the accompanying text-input field your desired end date. Always use the format shown, including a full four-digit year.
- View Calendar offers a pop-up window called "select a month." This includes a menu of 12 month names, a text-input field for the year (always use four digits), and a BUILD button to request calendar construction for the month and year that you specified. Using this calendar does NOT change the arguments for Start Date and End Date, however, which you must type in by hand in their own fields.

Output Options

offers three mutually exclusive buttons (Format, Tab, Plot) that control how the (bottom of page) Submit button performs, plus two other report-specification settings (right-hand side), as follows:

- Format causes the Submit button to generate your specified report and display it to your terminal screen as standard textual output from LRMUSAGE (hence, not the same as the like-named -f option).
- Tab (-t) causes the Submit button to generate your specified report with tabs and then offer a pop-up menu that lets you deposit it in your choice of local text file. This output is intended for subsequent input into another program, such as a database or spreadsheet.
- Plot causes the Submit button to generate your specified report and then plot it to your screen as a colored bar graph showing usage (y-axis) versus divisions, users, or banks (x-axis), depending on what values you requested.

Minutes	(-time) specifies (using a pop-up menu) the unit of time for your usage report (and for its reporting threshold, next item), where the mutually exclusive choices available are minutes (the default, whose name also identifies this otherwise unlabeled field), hours, days, and seconds.
0.0	is an unlabeled text-input field that specifies the time-reporting threshold for your usage report (only times exceeding this threshold are reported). The threshold units are minutes, unless you change them by using the pop-up menu described in the previous item.

HPSS Reports Interface

This OCF web interface is for authorized users only, not for general users.

Static Reports Web Interface

The web interface of Static Reports menus looks quite different than the web interface to [LRMUSAGE](#) (page 30) and to the [Job Status](#) (page ?) reports.

The Static Reports interface offers a jump-page list of preprocessed, monthly, comparative usage reports for selection and display. These reports are primarily of interest to system administrators and those monitoring the funding and consumption of computing resources (for example, ASCI Alliance partners). A table lists each month of the current year in chronological order and offers four report formats for that month (in columns) under these four headings:

Stacked Bar	visually compares the percentage of time used by each LLNL division on each machine and cluster (open interface shows only open machines, secure interface shows only secure machines). The plot is a labeled, color-coded bar chart.
Pie	visually compares the percentage of time used by each LLNL division on each machine and cluster (open interface shows only open machines, secure interface shows only secure machines). The plot is a labeled, color-coded pie chart covering the same data as the bar chart above.
Gflop	visually compares the peak gigaflops achieved on each machine and cluster (open interface shows only open machines, secure interface shows only secure machines). The plot is a labeled, color-coded bar chart.
User Usage	tabulates the percentage of available time used on each machine and cluster (open interface shows only open machines, secure interface shows only secure machines), where the output is organized first by machine, then alphabetically by bank and by user login name within each bank. Users consuming less than 1% of the available time are omitted.

Managing A Bank

Those few users with coordinator privileges can create or remove banks, assign or delete users in those banks, and spread resource shares among the banks and users that they manage. See the [introduction](#) (page 4) to this manual for a quick review of the LC accounting concepts and definitions relevant to bank management (with links to details elsewhere).

To see if you have coordinator privileges for a bank, type

```
bac -u uname
```

where *uname* is your LC login name. The BAC output reveals every bank to which you belong and marks with a C those for which you are a coordinator (you are also a coordinator for every child of every C-marked bank in the BAC report).

To manage a bank for which you have coordinator privileges, first log on to the LC machine where you want your banking changes to apply (otherwise they may apply to the wrong LCRM "resource partition"). Then execute LC's "Livermore resource manager" utility (formerly called "production control system manager" or PCSMGR) by typing

```
lrmngr
```

Respond to its lrmngr> prompt with the interactive commands shown here for the specific tasks that you want to perform.

CREATE A NEW BANK:

```
create bank bname parent pname share nn
```

creates a new bank called *bname* as a child of existing bank *pname* (for which you must have coordinator privileges) with resource shares *nn* (any integer, whose *ratio* to other bank shares is really more important than the absolute value).

```
show bank bname
```

reports the properties of bank *bname* to confirm that your previous step had the results that you intended (if not, you can alter the properties with the steps below).

ADD A USER TO A BANK:

```
permit user uname bank bname share nn [default]
```

assigns the user with login name *uname* to existing bank *bname*, gives the user *nn* resource shares, and optionally declares *bname* to be the user's default bank.

```
show user uname
```

reports the bank(s) and share(s) assigned to user *uname*, to confirm your changes regarding *uname*.

DELETE A USER FROM A BANK:

delete user *uname* bank *bname*

deletes the user with login name *uname* from bank *bname* and cancels that user's assigned shares.

CHANGE ASSIGNED SHARES (ALLOCATIONS):

update bank *bname* share *nn*

changes to *nn* (any integer) the shares assigned to existing bank *bname* (remember that relative allocations are more important than absolute ones).

update user *uname* bank *bname* share *nn*

changes to *nn* (any integer) the shares assigned from existing bank *bname* to the user with login name *uname*.

DELETE AN EMPTY BANK:

(first, delete every user as shown above, then...)

delete bank *bname*

removes from the bank hierarchy the existing bank *bname*. You must delete every user from that bank with DELETE USER (shown above) before you can remove the bank itself.

RENAME A BANK:

rename bank *oldname* *newname*

changes the name of an existing bank from *oldname* to *newname*, while leaving its users and share assignments unchanged.

You can also indirectly confirm the effects of using these LRMMGR commands by typing

pshare -T *bname*

to see the full current names and shares of the specified bank and all of its child banks (or use instead the -t option to see the login names and shares for all *currently active* users assigned to the specified bank or its children). Starting in February, 2006, LCRM makes no distinction between interactive banks and batch banks when it reports "time charged" by the jobs that it manages. Contact the LC hotline (925-422-4531, open e-mail lc-hotline@llnl.gov, secure e-mail lc-hotline@pop.scf.cln) for personal help in carrying out any of the bank management steps described here.

BRLIM Resource Limits Reporting Tool

BRLIM Execute Line

ROLE:

BRLIM reports on current "resource partition limits" (explained below) for users or banks. Like LRMUSAGE, BRLIM does *not* report on the resources used by any specific job, but rather on the total limits and resource commitments that affect all jobs for a user or a bank.

LCRM assigns every production machine to a "resource partition" (a cluster of similar machines managed together). And it assigns to every bank and to every user/bank combination a limit (perhaps 0 or infinite) on:

- the number of JOBS the user (or bank) can run at once in each resource partition,
- the number of NODES that can be committed to the user (or bank) at once in each resource partition, and
- the amount of NODE TIME (in minutes) that a user (or bank) can consume in each resource partition.

LCRM will not schedule a job on a machine if doing so would violate any of these limits for the resource partition to which the machine belongs (it follows that if BRLIM reports any limit as 0, no job for that user or bank will run).

REPORTING:

PSTAT reports jobs that cannot run because a resource partition limit has been reached by means of three special "status" values:

- | | |
|----------|--|
| JRESLIM | means that the job would exceed its partition's maximum job limit. |
| NRESLIM | means that the job would exceed its partition's maximum node limit. |
| NTRESLIM | means that the job would exceed its partition's maximum node-time limit. |

BRLIM reports the current aggregate LIMIT and USED values for each of these three variables (JOBS, NODES, and NODE TIME) for the user (and default bank) who runs it, or for other specific users or banks on request. Because these limits function somewhat like "allocations" in controlling which jobs run and when (or where) they run, BRLIM can help plan job flow or detect bottlenecks in job execution by LCRM (but BRLIM does not monitor individual jobs; for that use PSTAT). For more background on "resource partition limits," how LCRM uses them to manage batch jobs, available exemptions, and interactions with local limits, see the "Resource Partition Limits" section of LC's [LCRM \(DPCS\) Reference Manual](http://www.llnl.gov/LCdocs/dpcs). (URL: <http://www.llnl.gov/LCdocs/dpcs>)

MOAB MACHINES:

BRLIM does not run on nor report on any LC machines where jobs are scheduled by Moab rather than by LCRM. Moab does not enforce the JRESLIM, NRESLIM, and NTRESLIM limits and so PSTAT will never report a job in one of those states on a machine scheduled by Moab.

USAGE:

To run BRLIM (on LCRM-scheduled machines only), type

```
brlim [options]
```

BRLIM then reads the LCRM database and reports (as plain text) the current LIMIT and USED values for the user(s), bank(s), or machine(s) that your options specify (see next section). If run without options, BRLIM reports on your limits and commitments for your default bank for the machine where you execute BRLIM (then ends). There are no prompts; BRLIM ends automatically after every run. The output is in labeled text columns, and you can optionally suppress the column headers. Unless you (or your target users) are actually executing LCRM jobs at the moment when you run BRLIM, it will usually report zero values in all columns (even if your LRMUSAGE numbers show much recent activity).

BRLIM Options

User options:

- u *ulist*** reports limits and commitments for the users whose login names (not real names) are specified in the comma-delimited *ulist*. The order of names in *ulist* determines the order in which BRLIM reports values. (-u cannot be combined with -T, -t, or -r.)

Machine options:

- m *host*** reports limits and commitments for the resource partition of which *host* (the name of one specific LC production machine, such as THUNDER) is a member. Without -m, BRLIM reports (only) for the partition of which the machine on which you run BRLIM is a member. You *cannot* use a list of machine names for *host*, but you can request reports for all LCRM partitions at once by using * (an "escaped asterisk") as the value for *host*.

Bank options:

- b *blist*** reports limits and commitments for the LCRM banks whose names are specified in the comma-delimited *blist*. The order of names in *blist* determines the order in which BRLIM reports values. (-b cannot be combined with -T, -t, or -r.)
- T *bank*** reports limits and commitments for *bank* and for all of its child banks, down to the level specified by -l (or without -l, for all levels). (-T cannot be combined with -b or -u.)
- t *bank*** same as -T except that it also reports the individual user limits (if any) within *bank* and each child bank covered. (-t cannot be combined with -b or -u.)
- l *lev*** (lower case ell) specifies how many levels *lev* to display in the tree of child banks under the *bank* specified with option -T or -t, where *lev* must be an integer greater than or equal to 0. Use -l only with -T or -t.
- r *bank*** reports limits and commitments for *bank* and for each of its parent banks, up to and including its root bank. (-r cannot be combined with -b or -u.)

Display options:

- 0|O** (zero, upper case oh) prevents display of records for all users or banks that now have zero normalized share (i.e., that have no currently active jobs). This sometimes causes no records to display at all.
- H** issues reports without an explanatory header over each output column (default BRLIM output has the header).
- h** issues a brief help message summarizing BRLIM's command-line syntax.

BRLIM Examples

This section shows typical BRLIM execute lines and the corresponding program output, for both user-oriented and bank-oriented queries.

[1] BRLIM User Query

GOAL: To request the current limits and commitments for two specific users (other than the person running BRLIM) for the LCRM partition where BRLIM executes.

STRATEGY: Run BRLIM with the target user login names (must be comma separated) following the -u option. Note that the output order is not alphabetical, but instead is the same as the order of the names on the execute line. BRLIM does not run on nor report on LC machines scheduled by Moab.

User: brlim -u schach,dclark

Rtne:

USERNAME	BANKNAME	+---JOBS---		+---NODES---		+---NODE-TIME---		
		LIMIT	USED	LIMIT	USED	LIMIT	-	USED
schach	bdivp	none	- 5	none	- 5	none	-	nolim
schach	bdivi	none	- 0	none	- 0	none	-	0:00
schach	phydiv	none	- 0	none	- 0	none	-	0:00
schach	micphys	none	- 0	none	- 0	none	-	0:00
schach	peregrin	none	- 0	none	- 0	none	-	0:00
schach	micphysi	none	- 0	none	- 0	none	-	0:00
dclark	xdivi	none	- 0	none	- 0	none	-	0:00
dclark	xdivp	none	- 4	none	- 4	none	-	nolim

[2] BRLIM Bank Queries

GOAL: To request the current limits and commitments for a specific bank on a specific machine (not where BRLIM runs), or for all machines (all LCRM partitions).

STRATEGY: (1) Run BRLIM with the target bank following the -b option and the target machine name (here, thunder) following the -m option.
(2) Run BRLIM again and replace the name of a specific machine (after -m) with the special "escaped asterisk" pair of characters (*) to request coverage of all available LCRM partitions. The -m option does not accept lists of machines as an argument. BRLIM does not run on nor report on LC machines scheduled by Moab.

User: brlim -b bdivp -m thunder --- (1)

Rtne:

BANKNAME	PARENT	+	----	JOBS	----	+	----	NODES	----	+	----	NODE-TIME	----	+
				LIMIT-	USED			LIMIT-	USED			LIMIT	-	USED
bdivp	b			none	-			5				none	-	145:44

User: brlim -b bdivp -m * --- (2)

Rtne:

BANKNAME	PARENT	+	----	JOBS	----	+	----	NODES	----	+	----	NODE-TIME	----	+
				LIMIT-	USED			LIMIT-	USED			LIMIT	-	USED
RESOURCE PARTITION: alc_part														
bdivp	b			none	-			0				none	-	0:00
RESOURCE PARTITION: yana_part														
bdivp	b			none	-			1				none	-	nolim
RESOURCE PARTITION: thunder_part														
bdivp	b			none	-			5				none	-	145:09
RESOURCE PARTITION: up_part														
bdivp	b			none	-			0				none	-	0:00
RESOURCE PARTITION: zeus_part														
bdivp	b			none	-			1				none	-	0:07

UINFO Bank Membership Reporting Tool

You can discover which banks any user belongs to, or who belongs to any specified bank, by using the locally developed tool called UINFO. UNIFO (at /usr/local/bin/uinfo) reports basic background information about you or any other specified user. To run UNIFO type

```
uinfo [[user] uname | [group] gname | [bank] bname]
```

where UNIFO tries to guess if its argument is a user name (login name), a group, or a bank, but accepts your disambiguating option if offered as shown here. There are no prompts, and UINFO ends automatically after every run.

With no argument at all, UNIFO reports a 5-line help message summarizing its options and ends. With an argument, UNIFO reports as shown below (specific responses may vary from one LC platform to another if you belong to different groups or banks on different machines, for example):

[user] *uname* reports the same information as FINGER about the user whose login name is *uname* (real name, office telephone number, default shell, and home directory) plus the specified user's (numerical) UID, all online (not storage) groups to which the user belongs, and all banks to which the user belongs.

[group] *gname* reports the login name (not real name) of every user who belongs to the specified group *gname*, in alphabetical order. UINFO reports only online groups, not storage groups (see the "Using Storage Groups" section of [EZSTORAGE](http://www.llnl.gov/LCdocs/ezstorage) (URL: <http://www.llnl.gov/LCdocs/ezstorage>)).

[bank] *bname* reports the login name (not real name) of every user who belongs to the specified bank *bname*, in alphabetical order.

Because UNIFO reports on any user (and on any group or bank membership as well), it is more versatile than the similar LINUX tool USERINFO, which only reports on the user who runs it.

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.

(C) Copyright 2007 The Regents of the University of California. All rights reserved.

Keyword Index

To see an alphabetical list of keywords for this document, consult the next section (page 46).

Keyword	Description
<u>entire</u>	This entire document.
<u>title</u>	The name of this document.
<u>scope</u>	Topics covered in this document.
<u>availability</u>	Where these programs run.
<u>who</u>	Who to contact for assistance.
<u>introduction</u>	Definitions and background.
<u>tool-comparison</u>	Account, bank, usage tools compared.
<u>tool-chart</u>	Diagram compares key alloc. tools.
<u>account-tools</u>	ACC, NEWACCT, DEFACCT, etc. explained.
<u>bank-tools</u>	BAC, NEWBANK, BT, etc. explained.
<u>usage-tools</u>	LRMUSAGE, PSTAT, and PSHARE introduced.
<u>former-tools</u>	The tools replaced by LRMUSAGE.
<u>bankusage</u>	Obsolete tool replaced.
<u>userusage</u>	Obsolete tool replaced.
<u>padusage</u>	Obsolete tool replaced.
<u>lrusage</u>	Utility role and features explained.
<u>lrusage-execute-line</u>	How to run LRMUSAGE, pitfalls.
<u>lrusage-options</u>	Options grouped by task.
<u>lrusage-examples</u>	Sample LRMUSAGE reports.
<u>user-examples</u>	User-oriented LRMUSAGE reports.
<u>bank-examples</u>	Bank-oriented LRMUSAGE reports.
<u>division-examples</u>	Division-oriented LRMUSAGE reports.
<u>web-interface</u>	Web interface to LRMUSAGE functions.
<u>lrusage-menus</u>	Web-interface menus for LRMUSAGE options.
<u>hpss-menus</u>	Web-interface menus for special users only.
<u>report-menus</u>	Web-interface for static monthly reports.
<u>bank-management</u>	Privileged steps to make, stock a bank.
<u>brlim</u>	Resource limits reporting tool.
<u>brlim-execute-line</u>	How to run BRLIM, pitfalls.
<u>brlim-options</u>	Options grouped by task.
<u>brlim-examples</u>	Sample BRLIM reports.
<u>uinfo</u>	Bank membership reporting tool.
<u>index</u>	The structural index of keywords.
<u>a</u>	The alphabetical index of keywords.
<u>date</u>	The latest changes to this manual.
<u>revisions</u>	The complete revision history.

Alphabetical List of Keywords

Keyword -----	Description -----
<u>a</u>	The alphabetical index of keywords.
<u>account-tools</u>	ACC, NEWACCT, DEFACCT, etc. explained.
<u>availability</u>	Where these programs run.
<u>bank-examples</u>	Bank-oriented LRMUSAGE reports.
<u>bank-management</u>	Privileged steps to make, stock a bank.
<u>bank-tools</u>	BAC, NEWBANK, BT, etc. explained.
<u>bankusage</u>	Obsolete tool replaced.
<u>brlim</u>	Resource limits reporting tool.
<u>brlim-examples</u>	Sample BRLIM reports.
<u>brlim-execute-line</u>	How to run BRLIM, pitfalls.
<u>brlim-options</u>	Options grouped by task.
<u>date</u>	The latest changes to this manual.
<u>division-examples</u>	Division-oriented LRMUSAGE reports.
<u>entire</u>	This entire document.
<u>former-tools</u>	The tools replaced by LRMUSAGE.
<u>hpss-menus</u>	Web-interface menus for special users only.
<u>index</u>	The structural index of keywords.
<u>introduction</u>	Definitions and background.
<u>lrusage</u>	Utility role and features explained.
<u>lrusage-examples</u>	Sample LRMUSAGE reports.
<u>lrusage-execute-line</u>	How to run LRMUSAGE, pitfalls.
<u>lrusage-menus</u>	Web-interface menus for LRMUSAGE options.
<u>lrusage-options</u>	Options grouped by task.
<u>padusage</u>	Obsolete tool replaced.
<u>report-menus</u>	Web-interface for static monthly reports.
<u>revisions</u>	The complete revision history.
<u>scope</u>	Topics covered in this document.
<u>title</u>	The name of this document.
<u>tool-chart</u>	Diagram compares key alloc. tools.
<u>tool-comparison</u>	Account, bank, usage tools compared.
<u>uinfo</u>	Bank membership reporting tool.
<u>usage-tools</u>	LRMUSAGE, PSTAT, and PSHARE introduced.
<u>user-examples</u>	User-oriented LRMUSAGE reports.
<u>userusage</u>	Obsolete tool replaced.
<u>web-interface</u>	Web interface to LRMUSAGE functions.
<u>who</u>	Who to contact for assistance.

Date and Revisions

Revision Date -----	Keyword Affected -----	Description of Change -----
26Jul07	<u>usage-tools</u> <u>brlim-execute-line</u> <u>brlim-examples</u>	BRLIM not on Moab machines. Moab does not enforce BRLIM limits. BRLIM not on Moab machines. Partition report updated.
07Feb07	<u>user-examples</u> <u>bank-examples</u> <u>division-examples</u> <u>lrmusage-menus</u> <u>brlim-examples</u>	GPS replaced, details updated. ILX replaced, details updated. MCR replaced, details updated. Role of defunct machines clarified. Updated; defunct partitions noted.
15Mar06	<u>introduction</u> <u>usage-tools</u> <u>bank-examples</u> <u>bank-management</u>	PSTAT -f now reports AGUs. PSTAT -f now reports AGUs. No interactive/batch bank split now. No interactive/batch bank split now.
18Aug05	<u>user-examples</u> <u>bank-examples</u> <u>division-examples</u> <u>lrmusage-menus</u> <u>brlim-examples</u> <u>introduction</u> <u>tool-chart</u> <u>account-tools</u>	LRMUSAGE -h case updated. LRMUSAGE default case updated. LRMUSAGE -pad, -pad2 cases updated. Many details revised, expanded. Both cases updated. Accounts obsolete with LCRM 6.13. Most account tools deleted. Accounts obsolete, most tools defunct.
19Apr05	<u>tool-chart</u> <u>account-tools</u> <u>bank-tools</u> <u>usage-tools</u> <u>bank-management</u>	PHSTAT added. NEWACCT, DEFACCT fail on BG/L. NEWBANK, DEFBANK fail on BG/L. PHSTAT comparison added. DELETE USER option replaces DELPERM.
12Jul04	<u>lrmusage</u> <u>former-tools</u> <u>index</u> <u>web-interface</u>	LRMUSAGE replaces PCSUSAGE everywhere. LRMUSAGE replaces PCSUSAGE. New keywords (old retained too). URLs changed (again).
06May04	<u>pcsusage-examples</u> <u>web-interface</u> <u>hpss-menus</u>	Details and cases updated. New OCF and SCF URLs for service. New special-user report option.
13Nov03	<u>introduction</u> <u>tool-comparison</u> <u>bank-management</u> <u>pcsusage-options</u>	DPCS manual name changed. LRMMGR completely replaces PCSMGR. LRMMGR completely replaces PCSMGR.

		Resource, partition names updated.
03Jun03	<u>introduction</u> <u>tool-chart</u> <u>bank-management</u>	DPCS gradually becomes LCRM. PCSMGR also executable as LRMMGR. PCSMGR also executable as LRMMGR.
17Jan03	<u>brlim-execute-line</u>	JRESLIM, NRESLIM, NTRESLIM replace RESLIM.
06Jan03	<u>brlim-execute-line</u> <u>brlim-examples</u>	Zero limit (RESLIM) case noted. Output changed to clarify columns.
03Sep02	<u>tool-chart</u> <u>bank-tools</u> <u>uinfo</u> <u>index</u>	UINFO added. UINFO instructions added. New section for new tool. New keyword for new section.
16Apr02	<u>introduction</u> <u>tool-chart</u> <u>brlim-execute-line</u>	Limits as pseudoallocations. BRLIM added. Cross ref to details in DPCS Manual.
12Mar02	<u>brlim</u> <u>pcsusage-examples</u> <u>web-interface</u> <u>index</u>	New section for new tool. New clusters replace old ones. Menu options updated. New keyword for new section.
07Mar01	<u>index</u> <u>bank-management</u> <u>bank-tools</u>	New keyword for new section. New section for bank coordinators. Cross ref to new section added.
19Oct00	<u>index</u> <u>web-interface</u> <u>introduction</u> <u>pcsusage-execute-line</u>	New keywords for new subsections. New section on PCSUSAGE by web. Cross reference to new section. Cross reference to new section.
08May00	<u>index</u> <u>former-tools</u> <u>pcsusage</u> <u>usage-tools</u>	New keywords for new sections. Comparison of four tools added. New tool with options, examples. PCSUSAGE replaces older tools.
02Jun99	<u>bankusage-options</u> <u>userusage-options</u> <u>padusage-options</u>	Deleted Meiko as a partition. Deleted Meiko as a partition. Deleted Meiko as a partition.
02Sep98	entire	Links to DPCS Manual revalidated.
19Aug98	<u>scope</u> <u>introduction</u> <u>tool-chart</u> <u>usage-tools</u>	Elaborated, details added. Fair-share impact expanded. BT deleted, PSTAT added. Usage ambiguity stressed.

19Mar98	<u>who</u> <u>introduction</u> <u>tool-chart</u> <u>bank-tools</u> <u>usage-tools</u>	DOCGUIDE cross ref. added. Fair-share role explained. PSHARE added. BT now SCF only. PSHARE added.
---------	---	--

22Jan98	entire	First edition of this manual.
---------	--------	-------------------------------

TRG (26Jul07)

UCRL-WEB-200045

Privacy and Legal Notice (URL: <http://www.llnl.gov/disclaimer.html>)

TRG (26Jul07) Contact on OCF: lc-hotline@llnl.gov, on SCF: lc-hotline@pop.llnl.gov